REFERENCES
Linked references are available on JSTOR for this article:
https://www.jstor.org/stable/1805628?seq=1&cid=pdf-reference#references_tab_contents
You may need to log in to JSTOR to access the linked references.

Financial Risk and the Burdens of Contracts

By Herman B. Leonard and Richard J. Zeckhauser*

A young artist we know was recently approached by a mail-order marketing company affiliated with a major credit card. She was asked to develop designs for use on sets of high-quality china to be marketed nationally. The company initially offered to purchase the designs on a straight commission basis—our friend would receive a percentage of gross sales. She asked to be paid a fixed fee instead.

This request seems odd. It might indicate that she distrusted her artistic capacity for this task (she didn't), that she would slack off (she wouldn't), or that she distrusted their marketing skills (she did, but she hardly wanted to tell them). The company's royalty proposal provides a financial incentive for quality, but it conveys a mixed message. To induce a large investment of the artist's time, the company should have tried to persuade her that it was confident of selling many sets. The royalty offer suggests instead that the firm fears only a few sets will be sold, and wants to protect itself against excessive artist's fees per set.

The negotiation concluded with the company agreeing to pay an attractive fixed fee. This arrangement placed the marketing risk on the company, which is likely both to perceive it as smaller and to be better able to bear it no matter what its size. It also removed any incentives to undermarket that might have been created by a royalty agreement.

But the contract gives the artist no direct incentive to spend the extra time to make these designs superlative (such efforts would be repaid by attracting future business, however).

This contract clearly attends to financial risk. The effects of variance in the outcome are borne by the party better able to absorb them. But it is difficult to assess the risk-assignment aspects—how much our friend "paid" to get the insurance she bought, or for the negative signal about confidence she had to send to get it—because they are intertwined with a collection of other messages and complicated by other motivations. A basic tenet of economics is that joint gains from trade arise from differences between the parties (Leonard and James Sebenius, 1983). But the parties must simultaneously discover and exploit their differences; thus a contract becomes not merely a division of known gains, but a device for revealing them as well. Even relatively simple contracts are asked to carry a whole collection of informational and incentive burdens. Risk perceptions (Amos Tversky and Daniel Kahneman, 1974), risk preferences, and the chosen allocation of risk between the parties (Kenneth Arrow, 1971) are three elements of the collection; they may be difficult to see through the veil of the others.

I. Multiple Burdens

Contracts must carry multiple burdens when outcomes—and possibly inputs—are uncertain or unknown. Financial risk is recognized in the form of contracts: in what will be exchanged, under what contingencies. But this is not the only issue a given contract is designed to resolve, and often it is not the most important one. The form of contracts becomes particularly important when at least

one party is in doubt about outcomes because of uncertainty about the inputs that can or will be supplied by another.

In the traditional market paradigm, certainty and full information prevail, and the function of contracts is simply to divide the known outcome(s) efficiently. For example, wages must be set so that conditions of competitive equilibrium in labor markets are satisfied. Contracts simply specify exactly what will be exchanged, governed by what contingencies, and at precisely what price. Modern economics, in contrast, has examined a variety of forms of ignorance and the modifications of contracts that may arise to cope with them. Ignorance about random influences on *outcomes* (i.e., pure uncertainty about the future values of random variables) is in many ways the simplest form of missing information. Actual outcomes vary around their expected values; financial theory examines the consequences of that variance in a world where individuals have different perceptions of risk and different aversions to it. It prescribes modifications of contracts that efficiently divide both the expected outcome and the variance, and capture all relevant information possessed by at least a few participants in the market. If these were our only forms of ignorance, and if no information were closely held, financial markets could readily deal with risk, and we would be able to discern that they were doing so.

Two additional major forms of ignorance —about *characteristics* and about *performance*—have important consequences for contracting, however. Signaling theory is concerned with how we use observed purchases, which may include contract provisions offered, to reveal what cannot otherwise be discovered at such low cost (Michael Spence, 1974). Often we can arrange contracts or set conditions before entering contracts that will successfully distinguish people with materially different but unobservable characteristics, even when some of them would prefer to mask the differences. Agency theory starts from the proposition that principals cannot costlessly monitor the efforts or performance of agents they retain (Michael Jensen and William Meckling, 1976; John Pratt and Zeckhauser, 1985). Accord-

ingly, contracts must be concerned with providing motivation—that is, with enforcement (or at least support) of the intended bargain.

Simple contracts cannot always address all of these concerns simultaneously. The needs are at least distinct, and they are often in conflict. By analogy to a lesson of Jan Tinbergen, a few contract parameters may not be enough to address the concerns separately. Moreover, interactions of contract burdens can result in unexpected inefficiencies. For example, signals may be costly to obtain, but once purchased they do not typically affect marginal payoffs, and hence do not distort incentives. But when signaling is part of a complex package that must provide motivation and distribute risk, the resulting imperfect contract often does involve marginal distortions. Were there no risk aversion, such difficulties could be avoided, by merely letting the agent reap all the benefits from the outcome, that is, by selling him the company. If more than one agent is contributing effort, even this scheme will not work.

II. Worker Selection

These conflicts can arise even when there is no uncertainty about future random occurrences. Consider a firm that wants to hire the best available salesperson—the one who will sell most. Potential workers know exactly how much they would sell, but the firm does not. Each worker has the same reservation wage. Effort is not involved, just human capital (Gary Becker, 1964). The contract must both identify the best salesperson and divide the value of production between the firm and the worker. Clearly, no simple salary offer can distinguish among workers (either all or none would want the sales job, depending on whether the sales salary exceeded the reservation wage). A commission contract does provide this distinction. If the firm knows how good the best salesperson is, it can simply offer a commission rate low enough so that only the salesperson with greatest ability can earn more than the reservation wage. If the firm does not know the maximum ability in the population, it can still arrange the appropriate selection. For example, it can offer the job at auction,

hiring whoever agrees to accept the lowest commission rate.[1]

If we ask more of this contract, it too becomes insufficient. Suppose, for example, that actual sales are distributed around expected values known to workers, and that potential salespeople are risk averse while the firm is not. Then potential sales workers' desires to purchase insurance will conflict directly with the need to reveal which among them is the most able. The equilibrium must involve some use of royalties as a device to distinguish ability levels. But the participant selected through this sieve would prefer that the bargain be in the form of salary only. Similarly, a contract like this may partially reveal and accommodate differences in perceptions of risk, in work disincentives (and therefore prospective effort levels), and other features that both parties cannot perfectly observe—but only in exceptional circumstances will it be possible for a reasonably simple contract to achieve a fully efficient result.

A second example shows the interaction of risk aversion with other contract conditions more fully. Suppose we have three types of potential sales workers, each with reservation wage 100, with risk aversion and human capital (indexed by a normal distribution on output) as indicated:

| | Salesperson Type | | |
	1	2	3
Potential Sales Profit			
Mean	100	110	120
Standard Deviation	10	10	30
Risk Aversion	none	severe	severe

What form of contract will both select the right salesperson for the job and distribute risk efficiently between whoever is selected

and the risk-neutral firm? One strategy would be to offer the firm for sale to the highest bidder, at the price quoted by the second highest bidder. This has the standard incentive compatibility of any Theodore Groves (1973)/Edward Clarke (1971) scheme, but honest elicitation is not enough in this case. Given sufficient risk aversion on the part of bidders 2 and 3, the highest bidder will be 1—the least able of the lot. We can elicit signals about who is best by selling a stake in the firm—for example, we might set a commission rate of 10 percent and give the position to whoever will accept the lowest salary. But this may lead to the selection of bidder 2 over bidder 3.

A less common contract does manage to select the right bidder. Suppose we offer a salary of $99 and announce that we will also offer profit sharing to whoever bids the lowest share (we refer to such a contract as a "sliver," since a small fractional equity stake in the enterprise has been transferred to the agent). To maintain incentive compatibility, the sharing rate will be whatever is bid by the second lowest bidder. The fact that the fraction is small insures that the effect of risk aversion is removed, and the sliver system selects the most productive salesperson.[2]

It may not be practical to offer a profit-sharing plan of the sliver type—the commission rates may be so low as to seem silly, for example. Another contract, which we have not seen used, will work instead. Let the firm announce that it will pay $99 in salary plus a $2 bonus (an amount small enough not to incur risk-aversion losses) to whomever it hires if he or she meets a sales quota. The potential worker who volunteers to accept the highest quota will be hired, with the quota set at the level offered by the second highest bidder, for incentive compatibility.

[2] This contract breaks down as a selection device, however, if different bidders have different reservation wages. Suppose, for example, that bidder 3 has a reservation wage of 105. Given an excess productivity of 10, he or she is still the best candidate for the job. But now a contract of $99 (or any constant) plus a small equity stake will not be enough to separate bidder 3 from bidder 2.

The contract is a small bet, so risk aversion will not play a role. This contract based on a bonus for meeting a negotiated performance standard thus permits identification of the most able salesperson and appropriately allocates risk.

Both the need for a mechanism for revealing personal characteristics and the need to build in performance incentives are often in direct conflict with the parties' preferences for dividing risk. People generally know more about their own characteristics than others know about them, and they often have a better assessment of the risks involved (if any). Consider the entrepreneur-inventor approaching venture capitalists. The inventor may be unduly optimistic about the market's reception of new products, but almost surely has a better technical understanding of the potential—and the problems—of the proposed product. This more accurate assessment comes bundled with a conflict of interest about revealing it truthfully. Any contract with a supplier of capital must bear the burden of revelation as well as division. The most natural way to do this would be to offer contracts that will only be accepted by inventors with high subjective probabilities that their inventions will succeed—that is, contracts that will pay off handsomely only if the product is a considerable commercial success. But these are exactly the contracts that impose the greatest financial risk on the entrepreneur, the party to whom it is probably most costly.

If we resolve this conflict in favor of appropriate risk spreading—pay a fixed fee for the patent and put the inventor on salary—the contract can no longer reveal honest risk assessments by the most knowledgeable party. If we offer a strongly revealing contract—a small salary with a large royalty rate on sales beyond the breakeven point—the contract cannot allocate risk bearing efficiently. Moreover, the magnitude of the risk burden will vary directly with the incentive for the entrepreneur to work—when the acceptance of risk is being used as a signaling device, it often distorts work incentives at the margin. (This problem is not unique to signaling through risk, but in many situations the signal, once purchased, does not alter marginal incentives.) Neither extreme on this contracting spectrum is particularly attractive. In practice, we are likely to observe some form of compromise—which means that full optimality is sacrificed and that it may be hard to discern just how (and how much) the market arrangement has recognized financial risk.

III. Responses to Risk Sharing

Providing performance incentives in contracts also frequently militates against appropriate risk spreading. Consider the problem of rewarding managers. For a collection of well-understood reasons, shareholders want the firm not to be risk averse. But managers face career risk within the firm and are themselves understandably risk averse. As a result they may be unwilling to undertake sufficient risk. We often observe incentive contracts for managers (stock options, for example) that at least partially balance this (inefficient) caution. They are valuable only when the company does quite well—an event made more probable if the managers take greater risks. (A standard result of options theory is that ownership of options makes one risk preferring, or less risk averse.) But this performance incentive, deliberately designed to offset the risk aversion of managers, must proceed precisely by imposing greater risk on them. Indeed, studies of compensation for high-level executives have found very high levels of risk borne by the agents—evidently, the advantages from performance inducements and the selection gains from attracting risk takers compensate the losses from poor risk spreading. This makes it difficult to uncover the elements of these contracts that recognize and deal with financial risk—they are deeply intertwined with other burdens the contracts are carrying.

The need for incentives to balance the nonlinearity of rewards with performance can run in the other direction. Western legal tradition confers limited liability on shareholders, producing an asymmetry in payoffs around the point of bankruptcy. When a firm is near bankruptcy, it is protected from the full adverse consequences of any gamble, and this may induce risk-preferring behavior.

Indeed, it results in a fundamental managerial principle: "when your back is against the wall, roll the dice." Under these circumstances the firm may engage in gambles for which the risk-adjusted rate of return is below the risk-free rate, that is, it may accept lotteries involving a mean-sacrificing spread (*MSS*) of outcomes.[3]

IV. Robustness and Optimality

Since the many burdens contracts are asked to carry are often in conflict, full optimality is difficult to achieve.[4] But why do we not more often see more complex contracts, tailored to address a wider set of concerns? One answer, often advanced, is that customizing contracts is expensive. We propose an additional reason: like creatures too carefully adapted to a particular ecological niche, highly specialized contracts are vulnerable to small changes in conditions. For example, if we must address only risk aversion and selection on the basis of ability through a contract offer, a contract offering a high salary and a small equity stake in the outcome will suffice (the small equity stake attracts the more able; the fact that it is small means that it does not impose much risk and therefore involves no real loss due to risk aversion). If we need not worry about risk aversion, but must provide a marginal incentive for effort along with a selective mechanism to identify the more able, then a large equity stake in the outcome will generally be appropriate. In either case, a simple contract, specialized to address the precise conditions defining the contracting context, suffices.

Unfortunately, a small departure from the specified underlying conditions will make either contract seriously deficient. If agents are risk averse and we must both select the best workers and encourage effort at the margin, neither a small nor a large equity stake will work well—one provides too little incentive for effort, the other too much risk. Adding even a small payoff from effort to a contract context that originally involved only selection and risk aversion could dramatically alter the contract stipulations required to achieve optimality. The carefully arranged near optimality of a specialized contract crafted to address a given set of conditions can often be seriously upset by even a small dose of some other condition.[5] Thus familiar contracts, frequently simple in form, may not only be only cheaper to arrange, but may also be more robust. That is, they may achieve reasonable, though imperfect, results across a wider range of contracting conditions. This may be an important part of their attraction. Simplicity and robustness are, of course, not synonymous; we are merely observing that contracts designed to address fewer special features seem less likely to be blindsided by others.

The attractiveness of simple familiar contracts, in turn, makes it difficult to identify and measure the extent to which financial contracts appropriately recognize risk. Our ignorance is much more profound than simple market uncertainty. Contracts must specify divisions of benefits, but they must also

[3] Given a risk-free rate and a market premium for risk, a gamble is a mean-sacrificing spread if it involves below-market compensation for risk bearing. The expected return may be positive; it is still a sacrifice if the mean rises by less than we would expect, given the risk. Firms near bankruptcy might be induced to accept even gambles that involve an absolute reduction in mean in return for a sufficiently large spread. The propensity to engage in *MSS* can be offset by instruments such as bonds convertible to stock at the bondholders' option.

[4] In Appendix II (available from the authors), we describe a scheme for categorizing contracting situations and contract types, and comment on whether optimality can be achieved under various combinations of them.

[5] In many contracts, the degree of optimality obtained is a smooth function of the underlying conditions, and the characteristics of the optimal contract vary continuously as a function of the burdens being carried. For example, a small change in the risk aversion of the agent in a contract balancing performance incentives against risk spreading would call for only a small adjustment in contract terms. By contrast, when selection is involved a small alteration in conditions can lead to a switch in which candidate gets the job, and a consequent dramatic shift in the degree of optimality obtained. A major change in the form of the contract may be required to maintain optimality. For example, our negotiated quota arrangement can fail if the three salesmen have different reservation wages—and even a small change in reservation wages can take us across the boundary.

provide motivation even as they rely on it, and provide information about personal characteristics even as they make use of it. Because contracts must simultaneously address all of these needs, the result is either complicated arrangements or, more typically, uncomfortable compromises within simple contracts. This makes it difficult to see how any particular issue is being dealt with. We see only the congealed Gordian knot; the separate strands can no longer be distinguished.

REFERENCES

Akerlof, George, "The Market for Lemons," *Quarterly Journal of Economics*, August 1970, *84*, 488–500.

Arrow, Kenneth, *Essays in the Theory of Risk Bearing*, Chicago: Markam, 1971.

Becker, Gary, *Human Capital*, New York: Columbia University Press, 1964.

Clarke, Edward H., "Multipart Pricing of Public Goods," *Public Choice*, Fall 1971, *11*, 17–33.

Groves, Theodore, "Incentives in Teams," *Econometrica*, July 1973, *41*, 617–31.

Jensen, Michael and Meckling, William, "Theory of the Firm, Managerial Behavior, Agency Costs, and Ownership Structure," *Journal of Financial Economics*, October 1976, *3*, 305–600.

Leonard, Herman B. and Sebenius, James K., "Differences and Gains from Trade," mimeo., Kennedy School of Government, 1983.

Pratt, John W. and Zeckhauser, Richard J., *Principals and Agents: The Structure of Business*, Boston: Harvard Business School Press, 1985.

Spence, Michael, *Market Signalling*, Cambridge: Harvard University Press, 1974.

Tversky, Amos, and Kahneman, Daniel, "Judgement under Uncertainty: Heuristics and Biases," *Science*, September 1974, *185*, 1124–31.